VIEWPOINT

THE ROLE OF GOVERNMENT IN THE EVOLUTION OF THE INTERNET

Robert E. Kahn

From its origins as a U.S. government research project, the Internet has grown to become a major component of network infrastructure, linking millions of machines and tens of millions of users around the world. Very little of the current Internet is owned by, operated by or even controlled by governmental bodies. Most of the funding for Internet services comes directly from private sources, although the educational community in the U.S. receives most of its research funding from governmental sources. Increasingly, however, the provision of Internet communication services, regardless of use, is being handled by commercial firms on a profit-making basis.

This situation raises the question as to the proper long-term role for government in the continued evolution of the Internet. Is the Internet now in a form where government involvement should cease in its entirety so as to allow private sector interests to determine its future trajectory? Or is there still an important role for government to play in the future? This article concludes that there are a series of important contributions for government to make and a few areas where government involvement will continue to be essential to the long-term well-being of the Internet. In fact, as the Internet continues on its commercial and international trajectory, the role of government will continue to be essential, and may expand to include more involvement by governments around the world.

Internet Origins

The Internet originated as part of an ARPA research project on internetworking in the early-1970s. At that time, ARPA had demonstrated the viability of packet switching for computer-to-computer communication in its flagship network, the ARPANET, which linked several dozen sites and perhaps twice that number of computers into a national network for computer science research. Extensions of the packet-switching concept to satellite networks and to ground mobile radio networks were under-development by ARPA, and segments of industry (but notably not the traditional telecommunications sector) were showing great interest in providing commercial packet network services. At least three or four distinct computer networks were likely in existence by the mid-1970s and the ability to communicate between them was highly desirable if not essential.

In a well-known joint effort, circa 1973, Vinton Cerf (then at Stanford) and I (then at ARPA) collaborated on the design of an internetwork architecture that would allow packet networks of different kinds to interconnect and machines to communicate across the set of interconnected networks. The internetwork architecture was based on a protocol which came to be known as TCP/IP. During the period between 1974 and 1978, four successively refined versions of the protocol were implemented and tested by ARPA research contractors in academia and industry with the fourth version eventually being standardized.

The TCP/IP protocol was used initially to connect the ARPANET (based on 50kbps terrestrial lines), the Packet Radio Net (PRNET—based on dual rate 400/100kbps spread spectrum radios), and the Packet Satellite Net (SATNET—based on a 64kbps shared channel on Intelsat IV). The initial satellite earth stations were in the U.S. and U.K., but subsequently additional earth stations were activated in Norway, Germany and Italy. Several experimental PRNETs were connected including one in the San Francisco Bay area. At the time, there were no personal computers, workstations or local area networks available commercially and the machines involved were mainly large-scale scientific time-sharing systems. Remote access to time-sharing systems was made available by terminal access servers.

The technical tasks in constructing this initial ARPA Internet revolved mainly around the configuration of "gateways" (now known as routers) to connect different networks together and the development of TCP/IP software in the computers. These were both engineering intensive tasks which took considerable expertise to accomplish. Although industry would later offer commercial gateways and routers by making TCP/IP software available for some workstations, minicomputers and mainframes, initially these capabilities were unavailable. They had to be hand crafted by the engineers at each site by obtaining components from wherever they could be had.

In 1979, ARPA established a small Internet Configuration Control Board (ICCB) consisting primarily of members of the research community to help with this process and to work with the ARPA in evolving the Internet design. This process was also important in that it brought a wider segment of the research community more directly into the decision-making aspects of the Internet project which, until then, had been almost solely undertaken by ARPA. The ICCB was chaired by ARPA. The ICCB met several times a year and as the interest in the ARPA Internet grew, so did the interest in the work of the ICCB.

During this early period, the U.S. government (mainly ARPA) funded network R&D, supported the various networks in the ARPA Internet (they leased various components, bought others and contracted out the day-to-day operational management) and maintained responsibility for overall policy. In the mid- to late-1970s, experimental local area networks and experimental workstations developed by the research community were connected to the Internet by dint of engineering expertise at each site. In the early-1980s, commercial workstations and local area networks became available that were Internet compatible and the task of getting connected to the Internet was significantly eased.

In addition to R&D contracts, U.S. government contracts were issued for various infrastructure purposes such as to fund groups to maintain lists of hosts on the network and their addresses and to take responsibility for assignment and registration of host names and addresses. The government funded other groups to monitor and maintain the key gateways between its networks, in addition to support for the networks themselves. In 1980, the U.S. Department of Defense adopted the TCP/IP protocol as a standard. By the early-1980s, it was clear that the internetwork architecture that ARPA had created was a viable technology for wider use in defense.

Emergence of the Operational Internet

The Department of Defense had become convinced that for its use of networking to grow, it needed to split the ARPANET into two separate networks. One of these networks, to be known as MILNET, would be used for military purposes and mainly link military sites in the U.S. and perhaps around the world. The residual portion of the network would continue to bear the name ARPANET and continue to be used for research purposes. The use of the TCP/IP protocol would still allow computers on the MILNET to talk to computers on the new ARPANET, but the MILNET network nodes would be protected by virtue of locating them in protected sites, rather than in research labs and on university campuses. If problems should develop on the ARPANET, the MILNET could quickly be disconnected from it by simply unplugging the small number of gateways that connected them. In fact, these gateways were outfitted so that they could limit, if desired, the interactions between the two networks to the exchange of email.

By the early-1980s, the ARPA Internet was known simply as the Internet and the number of connections to it continued to grow. Recognizing the importance of networking to the larger computer science community, NSF began supporting a set of computer science researchers to connect to the emerging Internet via an effort called CSNET. This allowed new research sites to be placed on the ARPANET and be paid by NSF, and allowed new research sites to be connected via a commercial network (Telenet) which would be gatewayed to the ARPANET. A capability for supporting dial-up email connections was also provided. In addition, access to the ARPANET was informally extended to research colleagues at numerous research sites, thus helping to further diffuse the networking technology within the community. Also during this period, other federal agencies with computer-oriented research programs (notably DOE and NASA) created their own "community networks."

Although TCP/IP had been declared a defense standard a few years earlier, it was one of many standards (although the only one that dealt explicitly with internetworking of packet networks) and its use was not yet mandated on the ARPANET. However, on January 1, 1983, TCP/IP became the standard for the ARPANET, replacing the older host protocol known as NCP. This step was also necessary for the ARPANET-MILNET split which was to occur about a year later. With the use of TCP/IP now mandatory on the ARPANET, the rationale for connection of local area networks was evident and the growth of users and networks accelerated. It also led to a rethinking of the process that ARPA was using to manage the evolution of the network.

ARPA replaced the ICCB as a mechanism for dealing with network evolution with the Internet Activities Board (IAB). The IAB was to be constituted similarly to the old ICCB, but the many issues of network evolution were to be delegated to initially 10 task forces charted by the IAB and reporting to it. The chair of the IAB was selected from the research community supported by ARPA. Having previously assumed responsibility for overseeing the process that was used to evolve the Internet protocols, ARPA began to delegate the responsibility for certain aspects of the standards process to the IAB.

Following the CSNET effort, NSF and ARPA worked together to expand the number of users on the ARPANET, but they were constrained by the limitations placed by defense on the use of the network. By the mid-1980s, however, network connectivity had become sufficiently central to the workings of the computer science community that NSF became interested in broadening the use of networking to other scientific communities. The NSF supercomputer centers program provided a major stimulus to broaden the use of networks by providing limited access to the centers via the ARPANET. At about the same time, ARPA decided to phase out its network research program, only to reconsider this decision about a year later when the seeds for the subsequent High Performance Computer Initiative were planted by joint actions taken by the administration and then-Senator Gore.

In this period, NSF formulated a strategy to assume responsibility for the areas of leadership that ARPA had formerly held and planned for the fielding of an advanced network called NSFNET. NSFNET was to link the NSF supercomputer centers with very high-speed links, then 1.5Mbps and to provide for access from members of the U.S. academic community to the NSF supercomputer centers, as well as to one another. (For a brief period in the mid-1980s, there was a small initial NSFNET which linked the supercomputer centers with 64Kbps lines.)

In 1990, the last node of the ARPANET was decommissioned by ARPA. Replacing it was the NSFNET backbone and a series of regional networks funded in various ways, with most funded by the U.S. government or started with government funding and soon thereafter to be self-supporting. The effort by NSF greatly encouraged and expanded the involvement of many other groups in providing network services as well as using them. It followed as a direct result of the planning for the High Performance Computing Initiative which was being formed at the highest levels of government. Defense still retained the responsibility for control of the Internet name and address space, although it continued to contract out the operational aspects.

The Department of Energy (DOE) and the National Aeronautics and Space Administration (NASA) are heavy users of networking in the support of their respective missions. In the early-1980s, they built HEPNET (high energy physics net) and SPAN (space physics analysis net) based on the Digital Equipment Corporation's DECNET protocols. Later, DOE and NASA respectively developed ESNET (energy sciences net) and NSI (NASA science internet) to support both TCP/IP and DECNET services. These initiatives were early instrumental influences in the development of multiprotocol networking technology which was subsequently embraced by the Internet.

A key to the success of the Internet is to insure that the interested parties have a fair and equitable way of participating in its evolution.

International networking activity was also expanding in the early- and mid-1980s. Starting with a number of X.25 networks as well as international links to ARPANET, DECNET and SPAN, the networks began to incorporate open internetworking protocols. While initially the thrust was based on the use of OSI protocols, the same forces that drove the U.S. to use TCP/IP (such as its availability in commercial workstations and local area networks) also caused a growth in the use of TCP/IP internationally.

The number of task forces under the IAB continued to grow and, in 1989, the IAB consolidated its various task forces into two groups: the Internet Engineering Task Force (IETF) and the Internet Research Task Force (IRTF). The IETF was given responsibility for the "near-term" Internet developments. The IETF was responsible for generating options for consideration by the IAB as Internet standards. The IAB structure, with its task force mechanism, had opened up the possibility of getting broader involvement from the private sector without the need for government to pay directly for their participation. The U.S. government role continued to be limited to oversight and to provide financial support which facilitated the convening of the IETF in plenary session several times per year and sponsorship for many of the IRTF participants. By the early-1990s, even that mechanism needed to be augmented by the charging of nominal attendance fees to cover the out-of-pocket costs of the meetings.

The opening of the Internet to commercial usage was a significant development in the late-1980s. The initial step was to allow commercial email providers to use the NSFNET backbone to communicate with authorized users of the NSFNET and other federal research networks. Regional networks, initially established to serve the academic community, had taken on nonacademic customers as an additional revenue source in their efforts to become self-sufficient. NSF's Acceptable Use Policy, which restricted backbone usage to traffic within and in support of the academic community, together with the growing number of nonacademic Internet users, led to the formation of two privately funded and competing Internet carriers, both spin-offs from U.S. government programs. They were UUNET Technologies, spun off from a DOD-funded seismic research facility, and Performance Systems International (PSI) which was formed by a subset of the officers and directors of NYSERNET—the NSF-sponsored regional network in New York state and lower New England.

Beginning in 1990, the growth in the Internet (fueled significantly by the enormous growth on the NSFNET and including significant commercial and international growth) averaged upwards of 10% per month for months at a time. NSF also helped to stimulate this growth by funding incremental as well as fundamental improvements in Internet routing technology, as well as by encouraging the widespread distribution of network software from its supercomputer centers. Interconnections between commercial networks (and other networks) are handled in a variety of ways including the use of the Commercial Internet Exchange (CIX) which was established, in part, to facilitate packet exchanges among commercial service providers.

Recently, the NSF decided that further funding for the NSFNET backbone was no longer required. It also embarked on a path to make the NSF regional networks self-supporting over a several-year period. To assure the continued availability of network access to the scientific research community, NSF made competitively chosen awards to several parties to provide network access points (NAPs) in four cities. NSF also contracted with MCI for a very high-speed backbone service (initially at 155Mbps) linking the NAPs and several other sites. It also funded a routing arbiter to oversee certain aspects of traffic allocation in this new architecture.

The Internet Society was formed in 1992 as a private sector effort to help promote the evolution of the Internet. As a professional society, it has individual members but the bulk of its initial support comes from industrial contributions. In 1992, the IAB was reconstituted as the Internet Architecture Board, which became part of the Internet Society. It delegated its decision-making responsibility on Internet Standards to the leadership of the IETF, known as the Internet Engineering Steering Group (IESG). While not a part of the Internet Society, the IETF is the mixing bowl in which the candidates for future protocols emerge. The Internet Society now maintains the Internet standards process, and the work of the IETF is carried out under its auspices.

Issues for Consideration

As the Internet continues to grow, the roles of the various participants in developing and evolving standards naturally come under closer scrutiny. When the financial implications of Internet standard decisions were relatively small, the current standards process that emerged from the research community proved entirely satisfactory. As the financial impact of such decisions becomes increasingly significant, the nature of the standards process will continue to change accordingly to allow more direct industrial involvement in the decision-making process. How this will ultimately play out is yet unclear. However, the vitality of the current process derives from the broad involvement of the many communities that have a stake in the Internet. Unlike typical standards processes that form consensus decisions and then proceed to implement them, the Internet process works essentially in reverse by a kind of grass-roots mechanism. Internet standards are those that have been tried, found by actual trial and error to be desirable and the resulting standards developed as a result of widespread implementation. No better model of a standards process has yet emerged that is as dynamic and agile to allow more direct involvement by industry.

Further, with the widespread internationalization of the Internet, scores of countries now have fundamental interests in its evolution. Within the U.S., the Internet is seen in many quarters as the starting point for the National Information Infrastructure. Around the world, there is growing recognition that the set of NIIs (assuming each country commits to developing one) should be compatible with each other along some yet unknown dimensions. Who should take the lead in making this compatibility occur? Is this a role for the private sector, for governments acting together, or for some hybrid approach? There is clearly a role for government here, at least to provide oversight, support and guidance if not active participation.

Independent of the resolution of these issues, a further concern is the viability of any entity that has no individual or organization with overall responsibility for its evolution. It seems fair to say that many of the traditional carriers would prefer this capability to be provided as a turnkey service by them on an end-to-end basis. Industry surely has within itself the capacity to provide many of the necessary capabilities, but history shows that a government role was necessary to make it happen in the first place. What guarantees are there that the same degree of vitality in its future evolution will take place if market forces alone determine what new capabilities are added to the Internet? Furthermore, the Internet offers distinct possibilities for providing bypass mechanisms to conventional service offerings by the regulated carriers. These possibilities create the potential to make it more difficult for the regulated carriers to compete effectively in those areas and also make it hard for the relevant government regulatory bodies to ignore.

Finally, the carriers can only go so far in providing Internet services. Ultimately, the communication pathways must enter the users' machine, pass through layers of software and end up in applications programs. The computer industry along with the many vendors of computer-related equipment must play a role in determining how this segment of the Internet will evolve. The nature of innovation in technology almost surely guarantees that many new technological options will continue to be generated from many different sources and make their appearance throughout the Internet. Thus, it appears inherent that no single entity can possibly be in charge of the overall Internet. A key to the success of the Internet is to insure that the interested parties have a fair and equitable way of participating in its evolution, including participation in its also evolving standards process. A proper role for governments would be to oversee the process to make sure it remains fair and equitable and evolves so as to meet public needs across a wide spectrum.

An international infrastructure such as the Internet will ultimately require countries to deal with many of the details that are now taken for granted. For example, Internet names and addresses may take on additional legal meanings in the various countries as they rely on the Internet to a greater degree. Trademark of Internet names and addresses is only one aspect of concern. Contracts of all sorts may have Internet names and addresses embedded within them. Countries need to gain a sense of acceptance for names and addresses without necessarily assuming responsibility for day-to-day operation of this aspect of the Internet. Computer viruses know no national boundaries. If a major security attack should involve multiple countries, how will those countries work together to respond to such a challenge?

Finally, the ability to conduct network-based business between countries will require the resolution of many legal issues including intellectual property protection, the formalization of legal contracts on-line and the ability to deal with associated customs and trade-related matters. At its core, many of these legal issues require the use of encryption technology which has been perhaps the most closely held of all the network-oriented technologies. How can this kind of capability be made available in the international arena in ways that are acceptable to the national authorities? More generally, how can other controversial issues which are likely to arise in the future be effectively discussed and resolved? Various subsets of these kinds of problems have arisen in the context of other international public networks including the telephone network and are thus not unique or entirely new. As the Internet continues to grow, many of the approaches developed in that context may be directly or indirectly applicable to the Internet. Some combination of public and private sector involvement will likely be required to deal with these problems more generally.

It seems uncontested that governments have a fundamental role to play in the funding of advanced research and development which can push the frontiers of technology and knowledge. Often, this will involve the development and use of pilot projects or testbeds to try out ideas and concepts in actual working prototype environments. It also seems clear that governments must provide the necessary oversight to insure the standards process is fair and equitable. Governments must also take responsibility for helping to resolve problems which arise where independent decision making by multiple countries (e.g., in legal, security or regulatory matters) introduces further interworking problems. In the case of the U.S. national infrastructure development, the U.S. government must provide the leadership in many dimensions, including the removal of barriers where they inhibit and can be removed, the insertion of legal, security or regulatory mechanisms where the national interest so dictates, and the direct stimulation of public interest sectors that require and merit government assistance (for instance research, education or certain network aspects of public health, safety and universal access). Other nations may find similar incentives for government involvement as well.

Two final observations seem appropriate to make. First, it will be essential to separate the process by which standards are selected for the Internet from the process by which the options are generated in the first place. The current situation is almost ideal: standards are selected by a process akin to ratification only after independent design and implementation processes have produced viable working options. This separation needs to be maintained.

Second, the most important use of the Internet, and indeed the NII, will be to allow individuals to communicate with each other and to rapidly access the information they require or desire. In many cases, this information will consist of intellectual property of others. There is also the opportunity for every user of the Internet to become a potential provider of information services, thereby vastly increasing the amount of information available. How much of this information may be deemed valuable in a literary or business sense remains to be determined, but much of it may be important in many other contexts. It is essential that we sensitize individuals to the value of intellectual property, and the need to protect it accordingly; this will also have the side benefit of encouraging others to develop and make available intellectual property of their own. A combination of ethics, technology and law are needed to insure the effective development of this important aspect of the network in the future.

Conclusions

Over a span of some 20 years, the role of the U.S. government in the evolution of the Internet has gone from that of complete provider to involved participant. While the U.S. government took the lead in virtually every aspect of Internet in the early days, it currently plays a more limited role. It is now a major funder of network R&D, and it provides a significant oversight role in the evolution of the Internet. Within the U.S., it provides direct support or even control for several key aspects of the Internet operation, such as the registration of names and addresses and insuring adequate backbone connectivity; it continues to stimulate the architecture of the Internet along healthy new directions.

However, the role of the U.S. government in the Internet has been steadily declining for several years, particularly as the private sector and international interest in Internet has increased. In the final analysis, though, there is a major continuing set of roles and responsibilities for governmental bodies to undertake (both in the U.S. and around the world). Governments must architect the way in which different countries cooperate on various aspects of the Internet and its use and they must continue to provide the oversight role in its evolution, both nationally and internationally. Other governments can, but need not, play the leadership role that the U.S. government has traditionally done in the U.S. Without such a role in the U.S., however, it is doubtful the NII will become a reality on its own. Without government involvement on an international basis, it is unlikely that a global information infrastructure will emerge or that the Internet will evolve in as vital and dynamic a way as has heretofore been the case.

If one takes a long enough perspective, network and computer technology may be seen to still be in its infancy, and many of the current uses of the technology reflect past practices carried out more effectively in new environments. The real challenge to society will be for the public and private sectors to work together in harnessing the yet untapped potential of new and increasingly powerful technology in the network-based setting of the NII, and to nourish and incubate the powerful, even revolutionary new ideas that have not yet been uncovered.

Robert E. Kahn *is the president of the Corporation for National Research Initiatives.*